

# Introduction to the Argonne Training Program on Extreme-Scale Computing (ATPESC)

**Paul Messina**

*Director of Science*

*Argonne Leadership Computing Facility*

*Argonne National Laboratory*



# Outline

- **Welcome**
- **A few words about Argonne National Laboratory**
- **Motivation of the ATPESC**
- **The curriculum**
- **Logistics**
- **Follow-on opportunities**



## *Argonne's Mission: To Provide Science-based Solutions to Pressing Global Challenges*

Through discovery and transformational science and engineering...

World-leading  
hard x-ray  
sciences &  
sources

Discovery  
science for  
energy

Leadership  
computing and  
computational  
ecosystem

Fundamental  
physics and  
accelerator  
capabilities

Materials &  
systems  
engineering  
solutions

and through use-inspired science and engineering


Energy Storage

Sustainable  
Transportation

Nuclear Energy

Environmental  
Genomics

National  
Security





# *Major Scientific User Facilities*









# Motivation for the ATPESC

- With the challenges posed by the architecture and software environments of today's most powerful supercomputers, and even greater complexity on the horizon from next-generation and exascale systems, there is a critical need for specialized, in-depth training for the computational scientists poised to facilitate breakthrough science and engineering using these amazing resources.



# The DOE Leadership Computing Facility

- Collaborative, multi-lab, DOE/SC initiative ranked top national priority in *Facilities for the Future of Science: A Twenty-Year Outlook*.
- Mission: Provide the computational and data science resources required to solve the most important scientific & engineering problems in the world.
- Highly competitive user allocation program (INCITE, ALCC).
- Projects receive 100x more hours than at other generally available centers.
- LCF centers partner with users to enable science & engineering breakthroughs (Liaisons, Catalysts).



# Leadership Computing Facility systems

	Argonne LCF	Argonne LCF	Oak Ridge LCF
System	IBM Blue Gene/P	IBM Blue Gene/Q	Cray XK7
Name	Intrepid	Mira	Titan
Compute nodes	40,960	49,152	18,688
Node architecture	PowerPC, 4 cores	PowerPC, 16 cores	AMD Opteron, 16 cores NVIDIA K20x (Kepler) GPU
Processing Units	163,840 Cores	786,432 Cores	299,008 x86 Cores + 18,688 GPUs
Memory per node, (gigabytes)	2	16	32 + 6
Peak performance, (petaflops)	0.557	10	27





# 10 Petaflops Blue Bene/Q - Mira

- **Mira – BG/Q system**

- 49,152 nodes / 786,432 cores
- 786 TB of memory
- Peak flop rate: 10 PetaFLOPs
- 3,145,728 hardware threads

- **Vesta (T&D) - BG/Q system**

- 2,048 nodes / 32,768 cores
- 32 TB of memory
- Peak flop rate: 420 TF

- **Tukey – Nvidia system**

- 100 nodes / 1600 x86 cores/ 200 M2070 GPUs
- 6.4 TB x86 memory / 1.2 TB GPU memory
- Peak flop rate: 220 TF

- **Storage**

- Scratch: 28.8 PB raw capacity, 240 GB/s bw (GPFS)
- Home: 1.8 PB raw capacity, 45 GB/s bw (GPFS)
- Storage upgrade planned in 2015



# Notional Architecture Trends

Systems	2012	2017 +1/-0	2022 +1/-0
System peak	20 Peta	100-300 Peta	1 Exa
Power	10 MW	~15 MW	~20 MW
Node concurrency	12	O(100)	O(1k) or 10k
Total Node Interconnect BW	3.5 GB/s	100-200 GB/s 10:1 vs memory bandwidth 2:1 alternative	200-400GB/s (1:4 or 1:8 from memory BW)
System size (nodes)	18,700	50,000 or 500,000	O(100,000) or O(1M)
Total concurrency	225,000	O(100,000,000) * O(10)- O(50) to hide latency	O(billion) * O(10) to O(100) for latency hiding



# High-level view of curriculum

- Computer architectures, mathematical models and numerical algorithms
- Programming methodologies that are effective across a variety of today's supercomputers and that are expected to be applicable to exascale systems
- Multiple approaches on unifying concepts and levels of abstraction that provide migration paths and performance portability among current and future architectures
- Approaches to building community codes for HPC systems, and methodologies and tools relevant for Big Data applications





# Curriculum tracks/sessions and their leaders

- Architectures – Pete Beckman
- Programming models with emphasis on scalability and performance scalability – Rusty Lusk and Rajeev Thakur
- Mathematical algorithms and software -- Lois McInnes and Lori Diachin
- Software engineering in scientific computing – Katherine Riley and Anshu Dubey
- Performance and debugging tools – Kalyan Kumaran and Scott Parker
- Visualization – Mike Papka
- Data-intensive computing – Rob Ross and Rob Latham
- Community codes– Katherine Riley and Anshu Dubey
- Application case studies– Katherine Riley and Anshu Dubey
  
- The leaders for each topic will chair their sessions



# Dinner and lunch talks

- **Purpose: present additional topics that will probably be relevant to your research at some point in your career**
- **Eight dinner talks by live presenters**
- **During some lunches we will show videos of some of the talks that were presented at a recent symposium that celebrated thirty years of parallel computing at Argonne**





**Key people - *plus lecturers (see agenda) and many others too numerous to list***

## **Organizing and Program Committee**


- Pete Beckman, ANL
- Richard Coffey, ANL
- Kalyan Kumaran, ANL
- Rusty Lusk, ANL
- Paul Messina, ANL, Chair
- Mike Papka, ANL
- Katherine Riley, ANL
- Rajeev Thakur, ANL

## **Steering Committee**

- David Brown, LBNL
- Lori Diachin, LLNL
- Thom Dunning, University of Illinois
- Geoffrey Fox, Indiana University
- Jim Hack, ORNL
- Marc Snir, ANL
- Jeff Vetter, ORNL

## **Administration and Local Arrangements**

- Cheryl Zidel, ANL
- Ashley Boyle, ANL
- Ginny Doyle, ANL





# Yes, the ATPESC is an intense program

- **Many lectures every day, followed by evening hands-on sessions**
- **Ideally we would cover all topics in more depth but the result would be a six-week program**
  - But few people's schedules would allow them to participate
- **The scheduling of some talks is not ideal, due to lecturer's schedules**
  - E.g., Jim Demmel's and Jack Dongarra's lectures
- **Slides will be posted online as soon as available**





# Thank you, DOE Office of Advanced Scientific Computing Research (ASCR)

- **This training program was made possible by funding from the Research Division of the Advanced Scientific Computing Research program of the Department of Energy's (DOE) Office of Science**
- **The funding is for three years**
  - The training program will be offered again in 2014 and 2015
- **Help us improve the training program**
  - Track evaluations
  - Overall program evaluation
  - Conversations or emails to any of us



# Logistics (1)

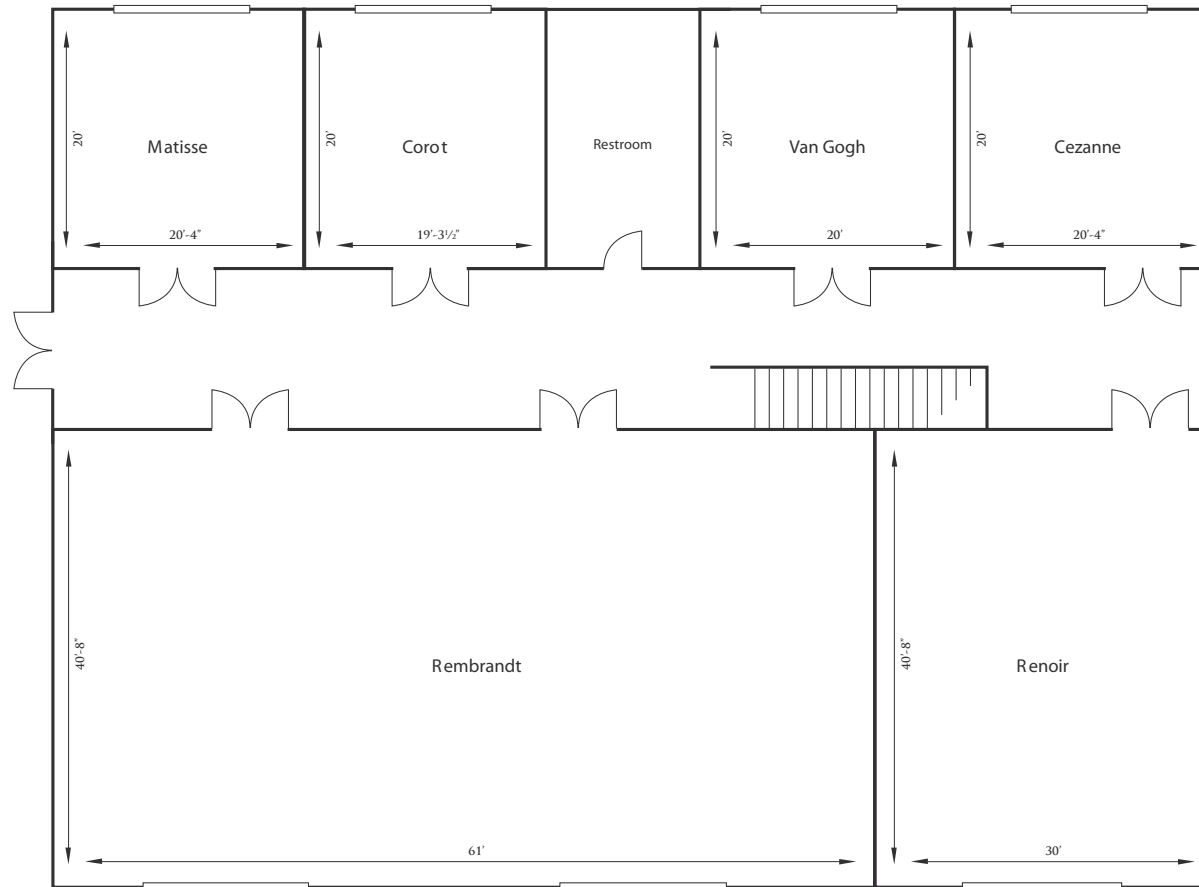
- All lectures and hands-on sessions in Rembrandt room
- All meals in Utrillo room on ground floor
  - Lunch and dinner presentations will be in this room
- Other rooms in ground floor might be used as needed
- Wi-fi SSID in this building is **Argonne**
- Password is **meeting2013**





# Diagram of Meeting Rooms: Second Floor

GALLERY HALL (SECOND FLOOR)

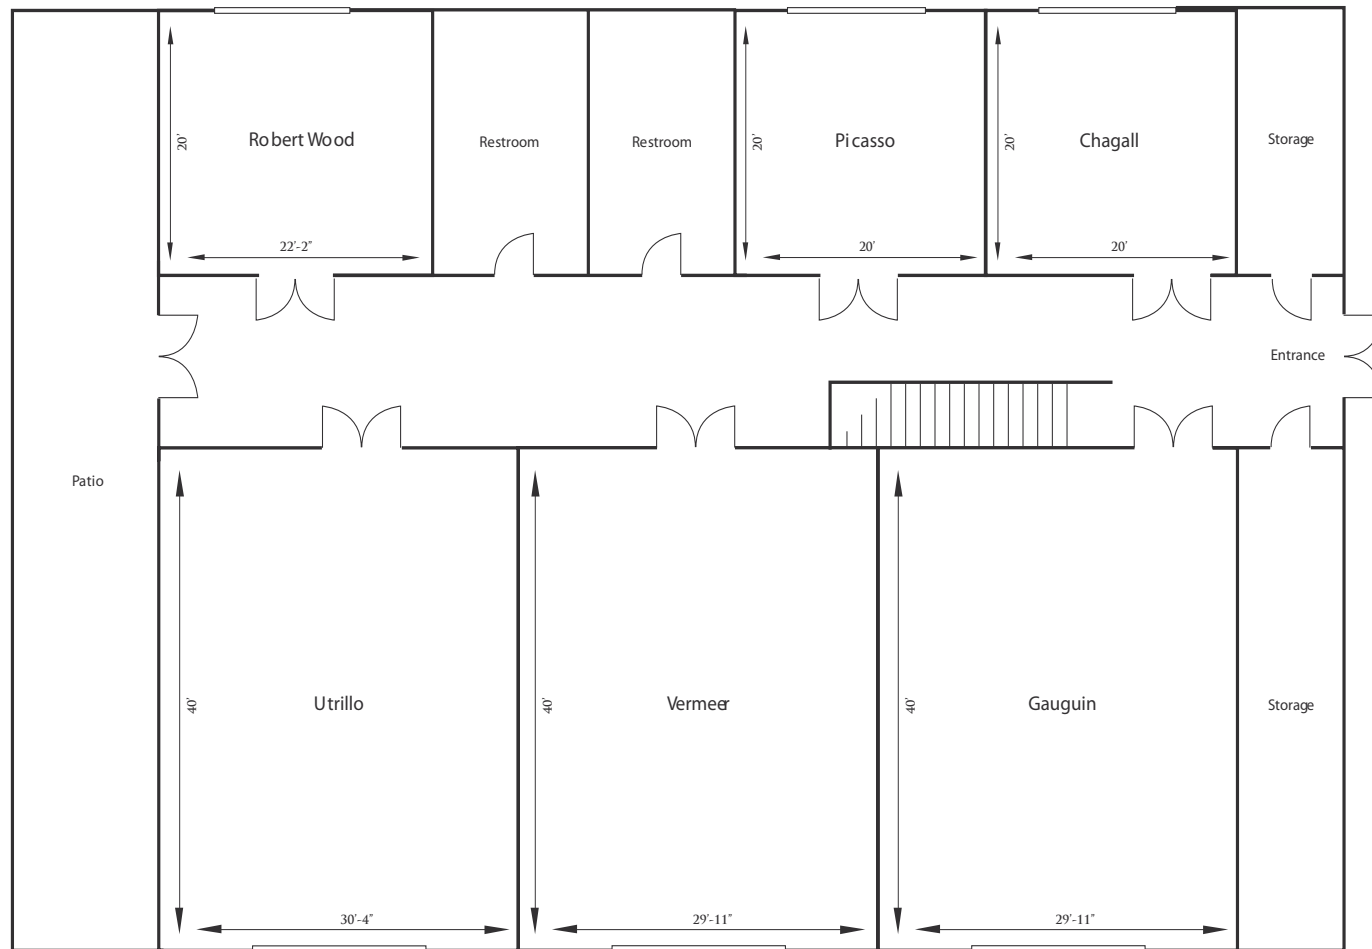


SCALE IN FEET  
0 5 10



# Diagram of meeting rooms: Ground floor

GALLERY HALL (FIRST FLOOR)



SCALE IN FEET  
0 5 10



## Logistics (2)

- All the lectures will be video recorded
- Recreational activities August 3 and 4: TBD



## Logistics (3) -- *Access to computing resources*

- **This evening after dinner, there will be presentations on how to access and use selected ALCF resources**
  - Vesta -- 2-rack Blue Gene/Q (one rack dedicated to ATPESC 24/7 tonight through August 9; both racks available during scheduled time slots if desirable)
  - Tukey -- visualization cluster with NVIDIA GPUs)
  - Mira (as time allows)
- **Training on using Keeneland and Titan will be later in the program**






# Whom to ask for help

- Local arrangements
  - Cheryl Zidel
  - Ashley
  - Ginny
  - (sometimes they will be in Picasso room on ground floor)
- **Computing issues**
  - Lalitha Mantrala
  - Robert Scott
  - Adam Scovel
  - Ray Loy
  - Hal Finkel
  - Others TBD



Some helpers may have a badge like this



# After the ATPESC

## Some opportunities

# Allocation Programs at the LCFs

	60%		30%		10%	
	INCITE		ALCC		Director's Discretionary	
Mission	High-risk, high-payoff science that requires LCF-scale resources*		High-risk, high-payoff science aligned with DOE mission		Strategic LCF goals	
Call	1x/year – (Closes June)		1x/year – (Closes February)		Rolling	
Duration	1-3 years, yearly renewal		1 year		3m,6m,1 year	
Typical Size	30 - 40 projects	50M - 500M core-hours/yr.	5 - 10 projects	10M – 300+M core-hours/yr.	100s of projects	.5M – 10M core-hours
Review Process	Scientific Peer-Review	Computational Readiness	Scientific Peer-Review	Computational Readiness	Strategic impact and feasibility	
Managed By	INCITE management committee (ALCF & OLCF)		DOE Office of Science		LCF management	
Readiness	High		Medium to High		Low to High	
Availability	Open to all scientific researchers and organizations <b>Capability &gt; 131,072 cores (16.7% of Mira)</b>					





# Educational and Job Opportunities @ ALCF

- **Research Efforts**

- Computational Science
- Computer Science
- Technical Communication
- Operations Research

- **ALCF Director's postdoctoral program (3 years)**

- **Divisional postdoctoral positions (3 years)**

- **PhD dissertation support**

- **Undergraduate and graduate internships**

- And advanced high-school level





For information on the educational and  
postdoctoral programs at Argonne  
National Laboratory

- <http://www.dep.anl.gov>



# Summary

- Thanks in advance to all of you for taking two weeks of your summer to participate in this program
- Questions?

